


MIT LIBRARIES



3 9080 00570434 8





WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

THE EMERGENCE OF A NEW TECHNOLOGY:
THE CASE OF NEURAL NETWORKS

by

Michael A. Rappa and Koenraad Debackere

June 1989

WP# 3031-89-BPS

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

THE EMERGENCE OF A NEW TECHNOLOGY:
THE CASE OF NEURAL NETWORKS

by

Michael A. Rappa and Koenraad Debackere

June 1989

WP# 3031-89-BPS

Michael Rappa is assistant professor of management, in the Management of Technology Group at the MIT Sloan School of Management. Koenraad Debackere is a visiting fellow at Sloan and a doctoral candidate in management at Ghent University in Belgium. The authors are particularly grateful to Susan Bjørner of the MIT Library for her expert assistance in electronic database searching, and to Professor Thomas Allen for his advice and encouragement. This research was sponsored, in part, by the Swedish Board of Technical Development.

MIT LIQUIDATION

THE EMERGENCE OF A NEW TECHNOLOGY:
THE CASE OF NEURAL NETWORKS

by

Michael A. Rappa and Koenraad Debackere
Massachusetts Institute of Technology

June 1989

Abstract

This paper examines the dynamic characteristics of R&D communities and their role in the emergence of new technologies. Using the case of neural networks as an example, it empirically investigates the growth and diffusion of the community using the literature published by neural network researchers as a source of information. The data indicate that the neural network community emerged with great speed, indicative of a bandwagon-effect, but only after experiencing many years of wavering enthusiasm and despair. The community survived intense controversy and lack of financial support, to persevere on the dedication of a small number of researchers who eventually succeeded in advancing their ideas to a state whereby they began to attract the commitment of others. As it grew rapidly, the community also became widely dispersed across hundreds of organizations in the public and private sector, and was held together by an elaborate grapevine that preserved the flow of information between laboratories.

Introduction

This paper explores the dynamic characteristics of R&D communities. The notion "R&D community" is used to signify a group of individuals, composed of scientists or engineers, who are committed to solving a set of interrelated scientific and technological problems, who may be organizationally and geographically dispersed, and who communicate in some way with each other.¹ The ultimate goal for some members of the community may be to create new knowledge, while for others it may be to apply existing knowledge in the creation of new products or processes. Furthermore, the community can include individuals employed in any type of organization, such as universities, private firms, new ventures, quasi-public corporations, and government research institutes wherever they may be located throughout the world.

This concept of an R&D community extends beyond traditional disciplinary and industry boundaries. Although it is common to speak of entities such as the "physics community," this is not what is intended by the term R&D community; nor is it meant to be a finer grain subset of researchers within a discipline, such as those involved in solid state physics. Rather, the community is normally interdisciplinary in nature and can include researchers from a wide variety of academic specialties. Moreover, within the private sector, it includes firms irrespective of their standard industrial classification. The only requirement is their focus on some part of the relevant set of problems. Thus, the community is defined here by the nature of the problem set and not necessarily by the end product, as is normally the case with an industry definition.

¹The notion of a research community has been developed extensively by scholars who have focused strictly on academic scientists and the so-called "invisible colleges" which arise in various problem areas (Ziman, 1984). The R&D community proposed in this paper seeks to extend the concept of an invisible college, to offer a more encompassing view, and in doing so, explore its relevance to technological development.

The dynamics underlying R&D communities are explored by studying the example of the neural networks community. This paper seeks to broaden the scope of current thinking on technological development which tends to emphasize the instrumental role of the firm in bringing new technology to the marketplace. Much academic research on the subject of technological development has been directed at understanding the functioning of firms (and industrial research laboratories, in particular), attempting to unravel the relationships between numerous organizational variables and success in developing new technologies. There is no doubt that previous work along this line of inquiry has yielded important insights about the management of technology development, but in focusing on the firm in isolation of the broader environment in which technological development occurs, it may have lost sight of some other critical dynamics in the emergence of new technologies.

In particular, it is proposed that technological development is not the exclusive domain of firms--or for that matter, a collection of firms in a given industry--but rather an activity which cuts across many types of public and private organizations. Furthermore, it is proposed that a new technological development may come about through the concerted efforts of a community of researchers which forms over time and that spans these diverse organizations. If this is indeed the case, then the progress of such communities may be better understood through a careful examination of how they function. This functioning is now further illustrated in the case of the neural networks community.²

Neural Networks: An Overview

A neural network is a type of information processing system modelled after biological processes, which has certain features that set it apart from the traditional and most widely used

²This study of neural networks is part of a larger study of the role of R&D communities in the emergence of new technologies.

type of information processing system, the digital computer based on the von Neumann architecture.³ The first distinctive characteristic of a neural network is that it is not programmed in the usual sense, but rather it is trained with data. This is an attractive feature because some information processing tasks are cumbersome from a programmer's perspective, and indeed devising the proper algorithm to perform a particularly complex task may be virtually impossible. It also implies that a neural network benefits from experience: as it processes more and more information while performing a task, it becomes increasingly more accurate in its response.

The second notable feature of a neural network is its degree of parallelism in processing a task. Unlike a traditional computer, with a single or small number of sophisticated central processing units, a neural network has a very large number of simple processing elements that operate simultaneously on a computational problem. The advantage of such massive parallelism is in its processing speed and fault tolerance (i.e., its ability to perform a task even though some processing elements fail.). In sum, although they may vary in form and function, neural networks, in essence, "*...are massively parallel interconnected networks of simple (usually adaptive) elements and their hierarchical organizations which are intended to interact with the objects of the real world in the same way as biological nervous systems do*" (Kohonen, 1988).

A clear distinction exists between neural computing and digital computing. Training a computer with data, or having it learn, are concepts quite foreign to the normal operation of the modern digital computer. The basic computer consists of a storage unit into which data are fed and a central processing unit with circuitry that can perform certain logic functions. It operates

³The neural network field is also referred to as "connectionism," "adaptive systems," or "neurocomputing." For a comprehensive overview of neural network technology, see the DARPA Neural Network Study (1988), or for a brief review see Hecht-Nielsen (1988).

in accordance with a stored program (or set of algorithms) that retrieves, manipulates, and stores bits of data (in binary form) in the course of its computational tasks. As Kohonen (1988) points out, biological neural systems do not apply principles of digital or logic circuits and therefore, the collective processes which are important in neural computing simply cannot be implemented by logic circuits. Therefore, unlike programmable digital computers, no machine instructions or control codes occur in neural computing. The circuits in neural networks do not implement recursive computation and are thus not algorithmic. Nor is the neural system's memory storage localized in a particular unit; instead, the memory function is distributed among the numerous processing elements. For this reason, Hecht-Nielsen (1988) claims that "neurocomputing is a fundamentally new and different information-processing paradigm--the first alternative to algorithmic programming."

The unique characteristics of neural networks may make them well-suited for certain classes of processing tasks that are otherwise difficult to perform using traditional computational techniques, and vice versa. For example, much like the human brain, neural networks are expected to perform best in applications requiring the processing of sensory information and controlling interactions with the environment, such as speech recognition, machine vision, robotic control, signal processing, and pattern recognition. But conversely, neural networks perform poorly in applications where rule-based problem solving can be implemented, such as handling mathematical calculations--something that traditional computers can do very well.

Neural networks have a long history of development. Early research has its roots in theoretical explanations of the brain and thinking processes proposed during the 1940s. During the initial decades, attention focused on the fundamentals of neural computing, including path breaking work on the formulation and elaboration of basic models to explain

such phenomena as adaptive stimulus-response relations in random networks. The 1960s were marked by the development of various implementations of neural networks, and in particular the single-layer "perceptron" by Rosenblatt. The perceptron was a watershed in neural network research, but for reasons which are discussed at greater length below, the work led to a schism among researchers generally interested in artificial intelligence. The controversy that ensued resulted in a decline in government support of neural network research in the United States after 1969 in favor of other AI techniques .

Nonetheless, work on neural networks continued during the 1970s, with a small number of dedicated researchers making significant progress on the theoretical foundations. The field then experienced a tremendous resurgence of interest in the 1980s, after significant developments had been made in both theory and application. Three general factors appear of paramount importance in explaining the recent growth of the field: (1) the evolution of the single-layer perceptron into a multi-layer system; (2) the rapid development of other technologies, such as semiconductors and computers, that enable researchers to develop, simulate, and diagnose neural nets of greater sophistication; and (3) significant progress in the understanding of neurobiological processes (DARPA, 1988). For example, building on advances in semiconductor integrated circuit technology, Sivilotti, Emerling, and Mead of the California Institute of Technology designed a VLSI collective decision circuit in 1985. This development and the subsequent designs of silicon models for neural computation set the stage for the design of neural network hardware. Even today, most neural nets are simulated on digital computers. However, the design of silicon models for neural computation brings the development of dedicated hardware one step nearer.

There are a number of neural networks that have been in the process of development during the last few years and several are now available in the marketplace. Among those neural

networks marketed commercially, some are software implementations that run on digital computers, others are hardware implementations (neurocomputers), and still others are a mixture of software and hardware. Neural networks are being explored for use in applications as diverse as Kanji character recognition, risk analysis in bank lending, and predicting process yields on a manufacturing line in real time. Although first generation products have attracted much attention, it is the promise of the next generation which is capturing the imagination of neural network researchers and potential users alike.

It is generally held that neural networks represent an emerging field of information processing that is a significant departure from established computer technology. Their potential seems enormous, though many problems still have to be solved before the technology reaches its full potential. The following sections of this paper examine the historical evolution of neural network community, seeking to shed light on the structural and behavioral characteristics of a rapidly emerging technology. Specifically, we empirically investigate three phenomena, which researchers themselves commonly speak of as the *bandwagon-effect*, *bootlegging*, and the *grapevine*. We also examine the controversy that arises between different communities, and the institutionalizing mechanisms used to preserve a community's momentum as it emerges. Given the experience of the neural network community, we will then discuss the general implications for a model for understanding the emergence of radically new technologies.

The Emergence of a New Technology

It is proposed that the emergence of a new technology is to a large extent the work of a dedicated R&D community. The community is formed and nurtured not by administrative decree, but by the autonomous actions of individual researchers who become intrigued by an idea and are committed to solving the problems necessary to make that idea work. In a sense,

it is a self-organizing process. Moreover, it is a subtle form of collective action in that it fosters cooperation among researchers without diminishing the competitiveness which naturally arises in pursuit of the rewards that will accrue to those who arrive at new knowledge and its application first. But even though researchers are competing with one another, they may nonetheless see themselves as members of a larger community from which they can draw on, and contribute to, in a variety of ways. As will be illustrated below, it is this combination of cooperation and competition that lies at the heart of R&D communities. This proposition is clearly illustrated in the case of the neural networks community.

Methodology

The data presented in this paper were obtained by studying the extensive body of scientific and technical literature generated by the neural network researchers themselves in the course of their work. Since communication among members is a defined characteristic of the R&D community, studying the documented, or formal, communication (in the form of papers and written presentations) among researchers is a convenient means for gaining insight into the functioning of the community. This literature provides us with a richness of background information about the technology itself, and subsequently about the behavior and the structure of the neural network community. In order to further systematize the study of the neural networks community, an electronic relational database was analyzed containing 2740 abstracts of neural network articles drawn from journals and conference proceedings since 1969. This analysis does not focus, however, on the well-established bibliometric methods for analyzing citation frequency and co-citation clusters. Instead, the abstracts were used in a different manner: namely, for the detailed information contained within each document.

Indeed, each abstract provides a wealth of reliable information about research activities within the community. The first advantage is its longitudinal character. By studying an R&D

community over a period of years, it becomes possible to show the dynamic patterns which emerge. The second advantage lays in the information obtained from the abstracts. Each abstract tells us who the researchers are and how their numbers vary from year-to-year. The abstracts also show what topics the researchers are working on and which organizations employ them. They also reveal the ties which develop over time among the different researchers and organizations by looking at the co-authorship of the papers and by showing the mobility of researchers between different organizations. By examining this information during a two-decade time span, one is able to visualize the structural changes in the R&D community, and also, to draw inferences about the behavior of the researchers dedicated to the particular technology.

Be that as it may, the literature does have certain non-trivial limitations. First, formal communication is relatively limited compared to the amount of informal communication that likely occurs among researchers at conferences and via telephone and computer networks. Partly as a consequence, individuals participating in the community may not become visible in the documented literature. Furthermore, the communication that does get published is slightly delayed in the process and not necessarily reflective of all the important details of the research. There may also be activity in developing the technology that is purposely hidden from public view by those who see it in their best interest to keep secret.

Nevertheless, given these limitations, the literature can prove to be remarkably useful in obtaining an initial understanding of the structural and behavioral dynamics of R&D communities. The analysis presented here seeks to limit the obvious deficiencies of using the documented literature by minimizing the sensitivity to publication frequency, by not attempting to ascribe more or less importance to a particular publication, and by analyzing the data historically. Thus, for example, we are not seeking to understand absolute magnitudes so

much as the dynamic trends over time, and we are not seeking to uncover technical secrets so much as obtaining a basic understanding of the people involved, the nature of their work, who they worked with, where they worked, and when they worked. Taken together, the data culled from the literature can provide a comprehensive picture of change over time within the community.

The growth of the neural networks community

An analysis of the neural networks database was conducted to determine the number of individual researchers worldwide, who were active each year in the research and development of the technology. If an individual is an author (or co-author) of a journal paper or presentation on the subject of neural networks in a given year, he is included as a member of the community; and he continues to be included as a member so long as he continues to be an author from year-to-year. In this way, membership in the community is not sensitive to number of publications by an author in a given year. This analysis yields a growth profile of the neural networks community between 1969 and 1988, which is shown in figure 1. The data indicate that until 1985, the community was rather small, with no more than 200 researchers active annually. However, after 1985 the community expanded very rapidly, from about 200 researchers to over 1000, which translates into a compound growth rate of over 60-percent. During this phase the community blossomed into a disciplinary collage composed of researchers with backgrounds in physics, computer science, electrical engineering, biological neuroscience, cognitive science, and linguistics, among others.

Such rapid expansion of a community is frequently explained by the *bandwagon-effect*; that is, the rush to join in on a trend, which has less to do with the underlying fundamentals than with the propensity for people to act on the basis of the actual or perceived actions of others. The bandwagon-effect contributes to an atmosphere of a race to the emergence of a




FIGURE 1: Growth of the Neural Networks R&D Community, 1969-88

new technology, a race for the rewards that accrue to those who are first to stake claims to new knowledge in the form of patents or papers. The result is the same regardless of the motivations of the individual researchers involved, whether it be fame, fortune or fad: a chain reaction that leads to an exponential growth of the community. The bandwagon-effect is significant because it underscores the large element of interdependence among researchers in their decision to participate in a technology's development. They undoubtedly influence each other. It is interesting to note that in the case of neural networks, we found the bandwagon-effect to occur similarly in both public and private sectors and in different countries.

An influential member of the neural networks community, Carver Mead, observes: "In a field like neural networks, one is usually too optimistic in the short run, but one is never optimistic enough in the long run" (DARPA, 1988). The bandwagon-effect may be fueled by an excessive degree of optimism among researchers, and may ultimately lead to the

proliferation of hyperbole in the form of unfounded and unrealistic expectations about the technology's potential. This is particularly true in conjunction with the mass media, where the neural networks community has received its fair share of "hype." The DARPA report addresses this issue directly, stating that the neural networks field lends itself readily to bold claims and sensational media headlines, such as proclamations of "thinking computers," and concludes that it is unlikely to dissipate: "Hyperbole's natural habitat is, after all, the territory between achievement and promise: the greater the potential of a new and mostly unexplored avenue of research, the greater the human tendency to imagine it opening pathways to both utopia and dystopia."

The proliferation of hype runs against the aspired norms of researchers, who as a group generally despise self-promoting hucksters who, in search for financial support, dare to oversell the potential of their research. For example, the founder of one neural networks company claims, "There is a whole lot of hype and nonsense in the field, which takes away from the credibility of the people who have been doing legitimate work. There has been important progress, but some people are grossly over portraying what the machines can do (Kinoshita and Palevsky, 1987)."

But, while it may well be contemptible to some, hype is a relative concept, and to a certain extent it may have the positive effect of attracting peoples' attention, including managers, venture capitalists, government officials--all of whom might have funds to allocate. Perhaps most important, is the need to attract other researchers to the field. Although neural network researchers may shrink when they see overly optimistic assessments of their field in print, they also understand the need for others (including students) to join in and become seriously committed to solving the difficult problems which confront them. The challenges may simply be too immense to be realistically accomplished in the near future unless the community attains

a certain level of participation, which researchers refer to as *critical mass*. For example, a leading academic researcher in neural networks and director of the DARPA study states, "I believe all the hype. What worries me is that many people don't realize how hard the problems are. Even if something works in the laboratory, it still takes ten years to get it out in the marketplace (Kinoshita and Palevsky, 1987)." The trick lies in generating a healthy optimism that creates enthusiasm rather than alienation among potential entrants to the community.

Cycles of Enthusiasm and Despair

The natural motivation of researchers to succeed in developing and applying new knowledge, in combination with a tendency toward optimism, yielded a very dramatic rise in the level of participation in the neural network community indicative of the bandwagon-effect. If the field is fertile and matches its promise, perhaps it will continue to enjoy a healthy expansion in participation. But what will happen if researchers discover, in the course of their experiments, that the problems they face are in fact much too difficult to resolve in the foreseeable future? Or that other technologies appear to be more promising? If researchers are so fast to enter a field, is it also likely that they might be quick to leave because their overly ambitious expectations cannot be met?

A closer examination of the community's growth profile in the pre-bandwagon phase reveals a pattern that is suggestive of the cycles of enthusiasm and despair that might befall a research community. Figure 2 shows the growth of the neural network community over the sixteen year period from 1970 to 1985, and makes a distinction between the core group of researchers (those who are usually active from year-to-year) and the rest of the community. Specifically, the core group is defined to include researchers who were actively participating in the community in at least three years of a given five-year period. Notice that the core group is relatively small (on average, between ten- to twenty-percent of the total community) and fairly

stable in size over the entire period, as should be expected. They are the researchers who keep the faith regardless of the obstacles that lie ahead, the true believers in the technology even when others disparage their work, the hopelessly committed. Success in their effort to push the field forward will preserve their place as founding fathers and garner the respect of their colleagues. However, sadly, if progress is slow or elusive, they will suffer a reputation as second-rate researchers with rather poor judgment.

Figure 2 also illustrates the cyclical pattern of participation among researchers whose commitment is less enduring than the core group. It is apparent from the data that the majority of researchers active in any given year are involved only briefly. Why this pattern arises is open to speculation. Perhaps it derives from the tendency for researchers to dabble in different areas which are tangential to their primary area of interest. It also might arise from the tradeoff researchers make in judging the likely potential of succeeding with one particular avenue of




FIGURE 2: Growth of the Neural Network R&D Community with Comparison of Core Group to the Total community

research versus another. Or it might to some extent occur as master's and doctoral students graduate and pursue other interests. In any case, the data suggest a degree of fluidity of participation in neural network community that stems from the movement of researchers between the variety of problem areas, given their expertise, they might pursue.

This fluidity of researchers into and out of the community is shown in greater detail in Figure 3. The database was partitioned into five, three-year periods in order to examine the extent of participation from period-to-period. In each period researchers were classified as new entrants (those not active in the previous period), sustainers (those also active in the previous period) and exits (those who were active in the last period by not in the present one). The degree of turnover from period to period is quite remarkable, with many researchers joining and others exiting each period.




FIGURE 3: The Flow of Neural Network Researchers into and out of the Community

An analysis of the extent to which neural network researchers who having left in one period return to the community in a later period was also conducted. It indicates that between 10- and 20-percent of entrants in a period (three through five) were previously involved in the community.

The foregoing analysis suggests that R&D communities, in their emergent stage, have a life of their own largely free of the formality that comes from organized programs of research, carefully budgeted and managed by a professional bureaucracy. Indeed, there is an air of informality and dynamism that seemingly derives from the independent decisions of researchers in search of promising new technological frontiers. How else could a community grow so rapidly or researchers move so fluidly, if not otherwise uninhibited by the red-tape usually associated with organizational decision-making in the public or private sector? But if researchers do act independently, how do they acquire the resources to sustain their effort?

The funding dilemma

The astounding growth of the neural networks community cannot be ascribed to a dramatic increase in public funding for neural network research. Although funding will certainly increase in the coming years (this appears from the intentions of the DARPA study), it was not likely the trigger for the dramatic increase in the number of researchers which occurred after 1985. Instead, the neural network community sustained itself and fostered its emergence not so much from formal public and private sector funding, as from dedicated researchers working more or less in the "shadows" of the laboratory. The terminology common among researchers is *bootlegging*, a word used to describe the practice of scrapping together resources to work, more or less surreptitiously, on interesting problems not sanctioned under existing budgets. It might take the form of working during lunch, evenings, or weekends, or taking advantage of the slack resources that typically exist in laboratories, or even using one's personal resources.

The practice of bootlegging is probably as old as the research laboratory itself, and seems to be a central element of every good story retold by researchers about their heroic effort to succeed in developing a radically new technology in spite of the powers that be. Bootlegging is a necessity--indeed, today it is even sanctioned by some laboratories--because of the very nature of research in new technologies. Important research is very difficult to judge in its early stages and thus, it is more likely to be rejected in favor of less risky, incremental work when formally evaluated. Researchers dedicated to a new and unorthodox technology are confronted with a difficult dilemma: they need more proof that their work will yield results before receiving resources, but without resources they are unable to do precisely that. Bootlegging enables fledgling research to go forward without the full knowledge and scrutiny of managers and other researchers, to a point at which the promise of the idea is clear.

To investigate the extent of bootlegging by neural network researchers, we analyzed the sponsorship of research presented at the First Annual International Neural Network Society meeting held in Boston in 1988. Of the 527 papers presented, 324 (61%) mentioned no sponsor at all. Only 19 papers (4%) reported research which was at least partly funded by the National Science Foundation. We found 43 papers (8%) that were at least partly funded by one of the US Defense Department agencies. Corporate involvement occurred in 108 instances (20%). One paper even described itself as "privately funded research," without any ties with the author's "current affiliations." One participant to the conference observes that in terms of travel money alone (perhaps a few million dollars), very little could be accounted for by funding agencies. Thus, although public research funds may start pouring into the neural network research laboratories, this kind of support clearly did not provide the trigger for the present interest in the field.

The relevant lesson from the practice of bootlegging is that researchers are driven by interesting problems, and not necessarily by what is looked upon favorably by funding agencies. Funding does not create a field, but rather allows it to flourish. If researchers make progress in a new area, the funding will hopefully follow suit. Whether it does indeed follow suit depends upon the researchers' fight for legitimacy among their peers, who may hold conflicting opinions about which research agendas are most promising.

The fight for legitimacy has been a particularly intense struggle in the case of neural networks. As stated previously, a relatively small number of dedicated researchers continued to work on neural network technology throughout the seventies--the so-called "wilderness years"--even though little funding was forthcoming and the technical feasibility of their concepts was heavily attacked by other researchers. At the heart of the controversy is the 1969 book, *Perceptrons: An Introduction to Computational Geometry*, written by two prominent MIT researchers, Marvin Minsky and Seymour Papert, who seriously questioned the potential of neural networks through a meticulous examination of the perceptron.

To Minsky and Papert, neural networks were (and still are) technically uninteresting, especially in comparison to the promise of rule-based AI techniques. But to neural network researchers, their treatise was seen as a blatantly underhanded attack to discredit neural network research in order to sway government funding away from the field. A leading neural network researcher recalls, "My impression was that Minsky and Papert defined the perceptron narrowly enough that it couldn't do anything interesting....It looked like an attempt to show that the perceptron was no good. It wasn't fair (Kinoshita and Palevsky, 1987)." In defense of their book, Papert admits:

Yes, there was *some* hostility in the energy behind the

research reported in *Perceptrons*, and there is *some* degree of annoyance at the way the new movement has developed; part of our drive came, as we quite plainly acknowledged in the book, from the fact that funding and research energy were being dissipated on what still appear to me (since the story of new, powerful network mechanisms is seriously exaggerated) to be misleading attempts to use connectionist methods in practical applications. But most of the motivation for *Perceptrons* came from more fundamental concerns, many of which cut cleanly across the division between networkers and programmers (Papert, 1988).

For his part, Minsky also admits that the attack on neural networks may have been too ardent, but defends his book claiming, "We were faced with a shortage of researchers in AI, and we were watching people throwing away their lives on perceptrons.... You would have to be crazy to be doing perceptrons when you could get so much more dramatic results with AI (Kinoshita and Palevsky, 1987)."

The bitter controversy lingers on and is perhaps as intense today as in 1969. (In an excerpt provided in Appendix A, Papert muses over the controversy in a recent and somewhat humorous allegory.) Although those within the neural network community contend that the limits posed in Minsky and Papert's book have by now been overcome, both authors remain skeptical about the future of neural nets: "Indeed, Minsky and I...suggest that the entire structure of recent connectionist theories might be built on quicksand: it is all based on toy-sized problems with no theoretical analysis to show that performance will be maintained when the models are scaled up to realistic size (Papert, 1988)." Indeed, Papert goes so far as to characterize the excitement among neural network researchers as buoyed by a "sustaining myth" that fits with the current intellectual fad to move away from rationalism toward more holistic ways of thinking.

In response to these attacks, neural network proponents have launched their own

vociferous criticisms of their detractors, suggesting that interest in neural networks stems, in part, from the disappointing results of AI research--this, despite the massive amount of funding the field has received over the years. Thus, in their view, it is the failure of AI that is breathing life into neural networks, and such diatribes as those of Minsky and Papert simply mask the declining promise of the established AI research agenda. However, in a more conciliatory tone, one commentator reflects on the controversy in this way:

What Minsky and Papert saw in the 1960s were researchers [who] were so enthusiastic about small successes that it blinded them to the real difficulties of the problems that they were trying to solve. Minsky and Papert saw understanding such important issues as representation, learning, and computational complexity in neural networks as fundamental to real progress in the field. Given the still shaky theoretical foundations of the field and the real similarities between the current wave of enthusiasm and that of the 1960s, its hard to fault Minsky and Papert for reminding us of these similarities (Will, 1988).

Despite the whirling controversy and despite the lack of funding for neural network research after the publication of *Perceptron*, some strong-minded researchers continued to believe in the potential of neural networks. Their diligence during lean years eventually paid off in the early eighties, as their ideas gained favor among a rapidly growing number of scientists and engineers employed across a diversity of disciplines and industries. The organizational and geographical spread of the neural networks community is examined in the following section.

Diffusion of participation and the grapevine

As described in the introductory section, R&D communities can extend beyond industrial boundaries, such that technological development is not wholly the domain of firms. R&D communities can consist of researchers employed by a variety of organizations, cutting across

university, industry, and government sectors. The neural networks case exemplifies this fact quite clearly.

For our purposes, the organizations involved in neural network research were classified as being either within the private or public sector. This second group is a mixture of university and government laboratories, although predominantly university. The distribution of researchers between the two sectors is shown in figure 4. Since 1969, a cumulative total of about 3800 man-years of research has been devoted to the development of neural networks. The fast majority of that effort took place in the public sector. Although industrial researchers had grown in numbers during the bandwagon-phase (reaching 18-percent in the last few years), prior to 1985 only between two- to seven-percent of all researchers were in the private sector.

The neural network community is an excellent example of how the efforts of academic researchers can contribute to the practical development of a new technology. Although neural networks have their roots in an academic setting, they have emerged in the eighties as a commercial activity. Many academic members of the community attempt to demonstrate the commercial viability of their research, and some even participate in private organizations in a variety of capacities. Like the more well known examples of semiconductor technology or genetic engineering, neural networks suggest that science and technology can and does become unified within R&D communities.

The neural network community is also quite international in scope, with more than a dozen countries participating. Indeed, researchers from relatively smaller industrial countries, such as Finland, have played an instrumental role in sustaining the technology's development through the seventies and now figure prominently in the field. Figure 5 shows




FIGURE 4: Sectoral Distribution of the Neural Network R&D Community

the international distribution of researchers among the six major industrial countries (the United Kingdom, West Germany, France, Italy, Japan, and the United States⁴) involved in the development of neural networks. In the post-1985 bandwagon phase, approximately seventy-to eighty-percent of the community resided in these six countries. It is interesting to note that the distribution among these countries varies over the entire period from 1969 to 1988, but that since 1985, the U.S. share has dominated the community, with about sixty-percent of the total. Thus, the rapid growth of the neural network community has been centered most strongly in the U.S.

As the neural network community grew in the period from 1982 onward, it became more widely diffused across public and private sector organizations. One way of viewing the extent

⁴According to the NSF, these six countries account for about 85-percent of the world's total research and development expenditures (NSF, 1989).




Figure 5: International Distribution of the Neural Networks R&D Community

of the community's diffusion is with a concentration ratio, similar to that used by economists, which measures the percentage of researchers employed in a set number of organizations with the largest research groups. For example, Figure 6 shows the change in the community's level of concentration for the top-5 organizations within the public and private sector from 1982 until 1988. Among neural network researchers in industry, the data indicate that prior to 1985, all were employed in just five firms. However, after 1985, as the bandwagon phase took hold, the level of concentration dropped rapidly, such that by 1987 only about 30-percent of industrial researchers were employed in the top-5 organizations. During this period, the number of industrial laboratories grew by more than a factor of ten, from five to over fifty.

In contrast, the level of concentration within the public sector portion of the neural network community declined gradually, from about 32-percent in 1982 to about 13-percent in 1988.




FIGURE 6: Diffusion of Neural Network R&D Community During Period of Rapid Growth

Although the public sector community had consistently been more widely diffused than the private sector, it nevertheless also experienced a similar increase in the number of participating university and government laboratories, from about 25 in 1982 to nearly 250 in 1988. Thus, in both sectors, we found a rapid diffusion of the community across different organizations.

Simultaneous with the growth and diffusion of the community, we see a tremendous increase in the number of firms that for the most part specialize in the development and application of neural networks. These firms are typically small in size, and in many cases they are ventures formed specifically for exploiting neural network technology. Investigation of the database led to the identification of 36 such firms. It appears that these firms form a bridge in the technology's transition from the public to the private sector. It is commonplace to find academic researchers deeply involved with the new ventures, taking part in their incorporation,

participating on the board of directors, acting as chief scientist or technical advisors, or channeling their graduates to positions as key members of the research staff.

The rapid growth and diffusion of the community likely has a beneficial effect on the commercialization of neural networks. More and more researchers, employed in an expanding array of organizations, seeking to develop and apply the technology in a increasing variety of ways, should almost certainly contribute further to the advances already made. As the community expands and spreads, it develops a powerful momentum that derives from the force of its numbers and the ingenuity of researchers independently working in laboratories in every corner of the globe.

However, the advantages of this expansion come with a cost: the diminishing ability of researchers to easily communicate with one another. This is a heavy cost, indeed, since it is the communication of information and knowledge among researchers that holds a community together. As a community grows in size, it quickly becomes a virtual impossibility for any given researcher to frequently communicate information to all others without some degree of coordination or structure. One formal means of communication, which will be discussed at greater length below, is through the published literature. However, this is a very rudimentary communication mechanism with severe limitations given the constant flow of new, complex information abounding from laboratories. Instead, another mechanism is needed: the *grapevine*.

The grapevine is an informal, but remarkably efficient network of researchers who facilitate the flow of information among different laboratories. Chances are the researchers are well acquainted with each other, perhaps having previously worked or studied together, or having become friends at a conference while commiserating over the years spent toiling over similar

problems. Moreover, the core researchers discussed above are likely central nodes in the grapevine. Indeed, over the two decades which have passed since 1969, a highly interconnected network among the different organizations involved in neural network research has developed. The network, illustrated in Figure 7, was analyzed by examining the database for instances where researchers in two different organizations were co-authors, or where a particular researcher (permanently or temporarily) changed his laboratory affiliation over the twenty-year period. The obvious assumption here is that researchers who collaborate on a paper are also likely to talk with one another even though they are not (or no longer) working in the same laboratory. In any case, this is possibly one of the most stringent criteria one can use to define a link between two organizations, such that the true grapevine is likely to be far more elaborate.

Notwithstanding this strict definition, a complex network of interconnected organizations appears that is likely representative of the core structure of the grapevine. The neural network grapevine, which in total connects 115 organizations, provides further supporting evidence for our basic proposition of the existence of R&D communities. Over the twenty year period, a total of 568 different organizations have been identified in the database. Thus, 20-percent of them are connected in the core grapevine. It is interesting to note that industrial research laboratories frequently appear in the central nodes. Moreover, while U.S. and European organizations appear fairly well integrated, Japanese organizations are less well connected to the core and some are part of a second, smaller network which is disconnected from the core.

Further analysis indicates that the complexity of the neural network grapevine has evolved over time and become particularly pronounced in the bandwagon phase of the community's maturation. In the pre-1985 era, very few organizations were linked together, and those that were constituted mostly of simple, bilateral linkages. It was not until after 1985 that the first




FIGURE 7: The Neural Network Interorganizational Grapevine

semblance of a more far-reaching grapevine became evident, and in the last year alone it has doubled in density.

So far, in analyzing the neural network community we have observed a number of dynamics that may hold special significance in understanding the emergence of a new technology. We have seen a community, which after years of wavering enthusiasm and despair, eventually take root, grow rapidly, and spread across hundreds of organizations in the public and private sector. While at the same time the community maintained its contiguity,

preserving the flow of information among its members via a rather extensive grapevine. But this remarkable surge of activity is not alone sufficient to ensure the success of the neural network community in commercializing its technology. The R&D community must continue to sustain a high level effort, or, in the vernacular of researchers, *preserve its momentum*. In the following section we describe some of the basic, and fairly obvious, institutions that evolve within R&D communities in order to facilitate their internal functioning.

Institutionalizing mechanisms

The preservation of a R&D community's momentum can be accomplished by a number of institutionalizing mechanisms, many of which can be seen in the neural networks case. The main purpose of these institutions is to ensure the community's healthy development in the post-bandwagon phase, by structuring the entry of new researchers into the field, easing the exchange of knowledge and information among members, and working to secure the funds necessary to sustain the effort long enough until the technology becomes commercially viable.

First of all, formal professional societies get established. Thus, the neural networks community created its "International Neural Network Society" (INNS), which now lists 3,500 members from 38 countries (with students comprising over one-fifth of the membership). The primary objective of such societies is to facilitate the transfer of information and knowledge among community members through organizing gatherings, such as conferences, workshops, and lecture series, and sponsoring the publication of a journal. In the words of its president, the Society works as an "interdisciplinary catalyst," with the goal of providing "information, education and a means of communication" across a range of disciplines in science and engineering (Widrow, no date). Toward this end, the INNS created its journal *Neural Networks* in 1988 and organizes an annual conference. (Another publication, the *Journal of Neural Network Computing: Technology, Design and Applications*, has been established

independent of the INNS.) The force behind these institutions tends to be the core group of researchers earlier identified, who serve as officers of the society and who sit on the editorial boards of the journals.

Beside allowing for the presentation of recent work, conferences provide an opportunity for the interpersonal exchanges that aid significantly in the solidification of collegial relationships--key to the functioning of the grapevine when researchers return to the laboratory. Between 1988 and early 1989, we were able to identify thirty-one conferences which were partly or totally devoted to neural networks technology. Perhaps the premier conference was the First Annual INNS Meeting in Boston in 1988, with over 600 papers presented by researchers from all corners of the neural network community.

Another mechanism serving to speed the communication process, are simple newsletters that are written to quickly disseminate the latest news on progress in the field. They can be in the form of a brief booklet (like the *Neural Network Review*), or even a "gossip sheet" (such as *Neurocomputers* and *Intelligence*), produced fairly inexpensively by freelancing researchers with a penchant for scuttlebutt. One variant on this is the firm-sponsored newsletter, such as Neuralware's *Connections*, which serves more as a means for publicizing products and recent applications.

Perhaps the speediest and most advanced mechanism for exchanging information is the communication network system that links laboratory computers across the world. In particular, the invention of the "interest group" or "bulletin board" format on these networks provides an astonishingly fast and efficient means for mass communication among members of the community. Although it is a electronic medium, like the telephone and the facsimile machine, the computer network interest group is more like the traditional printed publication in

its ability to reach large numbers of researchers simultaneously, only much more quickly. A neural networks interest group has already been in operation for at least a year, and the INNS intends to sponsor its own.

Another important institutionalizing mechanism, used primarily for the codification and transfer of the community's knowledge base, is the textbook. A number of neural network textbooks have appeared (among leading publishers is the MIT Press, with already nine textbooks in print). Textbooks also play a central role in the education of new recruits into the community, especially when woven into a coherent curriculum. This points to yet another mechanism, the formal graduate research program, which provides a direct means for channeling young researchers into the community and ensures that knowledge of the field will be handed-down from generation to generation. Therefore, it is not astonishing that leading universities have started offering graduate programs focused on neural networks. Already seven in number, these universities are highly central in the community grapevine. Along with the formation of graduate programs, specialized research centers and laboratories are also established. In this way, the community gains recognition as a separate entity within university settings, just as it does within the corporate world (through the creation of specialized new firms).

Lastly, there are structured efforts to ensure that funds will flow into the field in sufficient quantity to support researchers who are committed to developing the technology. The DARPA study, which was conducted with the help of more than fifty neural network researchers, is just one example of the community's effort to gain financial sponsorship. There have also been specific conferences and workshops dedicated to bringing together researchers with potential sponsors, in order to educate funding sources about the potential of neural networks and to educate researcher's about the intricate process of securing funds.

There is very little novelty to the institutionalizing mechanisms we have observed in the neural network community. Nonetheless, what is interesting is that these institutions emerged from the grassroots of the community, with remarkable speed and strength, and without central direction. It is truly a self-organizing process--the product of numerous researchers acting independently, but moving with a synchronism that derives from a shared goal. Taken together, these institutions are quite effective in preserving the community's momentum.

Discussion

Clearly, technological development is a complex process, but nevertheless the neural networks case suggests some basic elements that may contribute to a general theory of how the emergence of a new technology occurs. In this section, we offer a number of propositions that might fit within such a theory, given our observations of the structural and behavioral characteristics of the neural networks R&D community. It is readily admitted that this discussion is incomplete and that the propositions cannot adequately describe all of the subtleness of technological development as it might actually unfold. Rather, our intent here is to capture the essential "internal" dynamics--the functioning of R&D communities--that might become one component of a more comprehensive theory of emerging technologies.

First, we propose the theory begin with the assumption that technology is, in essence, a body of knowledge.⁵ Even though the ultimate goal may be to produce something, the currency of R&D communities is not so much actual things as it is the ideas, or theories, about how and why things work the way they do. Therefore, technological development can be

⁵The idea that technology is essentially knowledge has gained acceptance among scholars in several disciplines. For example, see Arrow (1962), Layton (1974), and Constant (1980). Constant, in particular, has pioneered the effort to understand technological communities with his study of the turbojet. The work of Katz and Allen (1985) focuses on technology development as an information processing and problem-solving activity. The intent here is to link the concepts of information, knowledge, and problem-solving into a unified model of technological development.

understood as an intellectual process that evolves over time, whereby new knowledge is created and applied in order to construct a new product or process. Of course, emerging technologies are not developed from scratch but are a combination of newly created knowledge and existing knowledge drawn from other epistemic realms. However, in the early stage of a technology's emergence, this body of knowledge is incomplete--one simply does not know everything one needs to know in order to make the technology work. The areas in which knowledge is lacking can be characteristically viewed as problems. If the technology is to be successfully reduced to practice, then new knowledge, in the form of solutions to the problems, will have to be found.

The central actors in this process are the individual researchers who become dedicated to solving the problems, and it is they who set the process in motion with their efforts to create and apply knowledge. In the course of their work, researchers perform three basic activities: (a) they produce information, (b) they transform information into knowledge, or in other words, they solve problems, and (c) they communicate information and knowledge to each other. We propose that the rate of progress in a technology's emergence is a function of how quickly problems are solved, which, in turn, depends on the amount of information produced, the number of diverse solutions attempted, and the extent to which information and knowledge is communicated among researchers. We expect that the more information available to a researcher, the more likely he is to arrive at a useful solution. Moreover, we anticipate that the more diversity in the types of solutions attempted, the more likely that critical solutions will be found. Lastly, we hypothesize that communication between researchers enhances the probability of finding useful solutions.

It is reasonable to assume that researchers are rational, in the economic sense that they are motivated by self-interest: that is, they are eager to solve problems because there are rewards

for those who do. The researcher's objective is to maximize the amount of knowledge he produces and can lay claim to before other researchers, because these claims have potential value. The actual value of a claim is contingent upon whether or not all problems necessary for reaching commercialization are resolved and the length of time taken. Expressed in another way, for a given piece of knowledge, it is more valuable in use, and the sooner it is used the more value it will have. A researcher need not produce all of the knowledge required to commercialize a technology, as long as his own knowledge claims are secured.

Since his choice is guided by the objective to maximize knowledge claims, the researcher is motivated to participate in those communities, within the realm of his disciplinary expertise, in which he believes there is a good probability of finding solutions that yield him valuable claims. If the actual probability is equal to or greater than expected (the task of solving problems is easier than anticipated), then the researcher will likely remain in the field; however, if it is lower than expected (solving problems is harder than anticipated), then the researcher might switch to another community with a higher perceived probability of success. The switching behavior of a researcher is moderated by the length of time spent contributing to a technology's development and the amount of knowledge claims accumulated. A researcher's specialized knowledge becomes a sunk cost: the more knowledge specific to a technology the researcher accumulates over time, the less likely he is to exit a community prior to reaping the rewards that come with reaching commercialization.

Furthermore, the researcher's decision to enter or exit a community is not a one-time choice, but rather is subject to frequent reconsideration based on new information. The attractiveness of developing a particular technology changes over time, as more information is produced and researchers reevaluate the probabilities of successfully solving the problems they face. Its attractiveness also may be influenced by changes in other technologies, and indeed a

variety of other events, all of which will be reflected in the researcher's decision.

A primary activity of researchers is to produce information and transform it into knowledge. By "information production" we mean the collection of new data through observation and experimentation. Knowledge is a distinct entity from information in that knowledge enables the researcher to do something (know-how) or explain something (know-why). Having information does not necessarily imply either.⁶ The researcher first must make sense of the information available to him--the cognitive process of transforming information into knowledge--in order to solve problems.

The processes of information production and transformation are time consuming; therefore, any individual researcher can only accomplish a certain amount of effort in a given period of time. The information he produces and the solutions he pursues are an expression of the researcher's own judgment and creativity, although he may be influenced by people with whom he collaborates. Unfortunately, not all attempted solutions will work satisfactorily. The probability that a particular solution will work successfully can only be subjectively determined *a priori*, and this is likely to change over time as researchers generate new information in attempting to implement solutions. Moreover, the relevance of any bit of information (in that it might contribute to the successful solution of a problem) cannot be determined *a priori*. For this reason, the value of information is less certain than that of knowledge.

The amount of information and knowledge available to a researcher depends upon how

⁶For example, a researcher runs an experiment and obtains negative results. He cannot use the information in the practical sense that enables him to do or explain something--that is, it is not knowledge. Nevertheless, the information may be instructive and eventually prove key in creating knowledge. Moreover, even though the information is not valuable in that it should be published or patented, it does have value in the sense that one who is aware of it will not waste time discovering it again.

much he can produce himself, or receive in the process of communicating with others. The communication of information and knowledge implies that a researcher can also gather information and knowledge produced by another researcher, and disseminate to others that which he produces or learns. For the most part, information is communicated informally by means of interpersonal conversations, whereas knowledge is communicated in the form of documented claims, such as with the submission of papers to refereed journals or patent applications.

The communication of information among researchers is influenced by the existence of organizational boundaries between researchers and their (and their organization's) economic interests. Organizational boundaries are important for two reasons: first they give rise to information asymmetries among researchers because they impede the flow of information and increase the cost of information gathering. As a result, organizational boundaries can slow the rate of production of knowledge within a community by reducing the amount of information available to each researcher. However, organizational boundaries can also enhance knowledge production to the extent it increases the diversity of problem solutions pursued by the community as a whole, since researchers in different organizations are likely to have different information sets.

To illustrate this point, assume the extreme conditions. In the first case, as the research community grows, all researchers are employed by the same organization. Thus, all researchers are exposed to the same information set, but the diversity of solutions performed is limited by the researchers' mutual influence. In the second case, as the community grows, each researcher is employed in a separate organization. Thus, each researcher is working from a different, limited set of information and performing independent solutions, such that the diversity of attempted solutions is maximized. Simply stated, in one situation each researcher

in the community has a wealth of information but a limited number of approaches in solving the problems being confronted; in the other situation, although the community can generate the same amount of information, the amount available to any particular researcher is small and the variety of approaches taken is great.

The communication impedance effect of organizational boundaries is overcome via researchers who act as technological gatekeepers--that is, researchers who tend to communicate with others in different organizations (Allen, 1979). Given the economically rational behavior of researchers, the communication of information across boundaries likely occurs as a form of quid pro quo (von Hippel, 1988). Furthermore, information is more likely to be the trading object of grapevines than is knowledge, because the value of information is indeterminate and because knowledge (which does have potential value) needs to be formally documented when disclosed in order to secure its value for the inventor.

The propositions presented in this discussion, taken together, suggest that R&D communities are freely functioning, loosely connected coalitions of researchers, who in their search for promising opportunities to expand on the frontier of knowledge, become dedicated to developing a common set of ideas. It is a self-organizing process, which grows out of the efforts of researchers who endeavor to solve interesting problems as they lock themselves into a struggle with nature. As researchers probe the frontier, sometimes they may find it stubborn and unyielding, and in other instances they may uncover an area which gives way to rich possibilities for the development of a new technology.

The speed at which a promising technology develops depends on how rapidly a community forms and grows, and the structural and behavioral characteristics the community adopts as it matures. Some characteristics promote cooperation between researchers, while others foster

competition. In combination, both elements are important contributors to the swift conduct of the problem-solving process: the competition that breeds with the rapid growth and diffusion of the community, escalating the amount of information produced and increasing the diversity of solutions pursued; and the cooperation which flows from the grapevine, enabling researchers to exchange notes on the latest experiments and to pursue solutions without squandering their precious time on known lost causes.

Conclusion

The objective of this research is to draw attention to the notion that R&D communities exist and that they may play an instrumental role in the emergence of new technologies. Although the importance of *academic* communities in the operation of science is generally accepted, very little is understood about how communities may contribute to technological progress. This continues to be the case, despite the fact that it is fairly common to hear people talk of entities, such as the "software development community," which are so intimately involved with technological development. What does it mean to be a member of a community? How do communities function? Can their internal operation tell us anything about the rate at which researchers are succeeding in solving the problems they face?

Using the development of neural networks as an example, we have investigated a few of the structural and behavioral characteristics of an R&D community. Specifically, we empirically assessed the growth and diffusion of the community using the literature published by neural networks researchers as a source of information. The data indicate that the neural network community emerged with great speed, indicative of a bandwagon-effect, but only after experiencing many years of wavering enthusiasm and despair. The community survived intense controversy and lack of financial support, to persevere on the dedication of a small number of researchers who eventually succeeded in advancing their ideas to a state whereby

they began to attract the commitment of others. As it grew rapidly, the community also became widely dispersed across hundreds of organizations in the public and private sector, and was held together by a rather elaborate grapevine that preserved the flow of information between laboratories.

It is unwise to read too much into the present finding without the careful examination of other R&D communities. However, the emergence of neural networks is suggestive of a self-organizing process that arises from the grassroots. It is hard to believe, given the evidence, that the direction of research can be centrally planned. The emergence of a new technology appears very much to be the result of independent forces, which become unified by a common commitment to a set of ideas and to solving the problems necessary to make these ideas work. It also suggests that technological progress might not be accelerated simply by brute force application of manpower alone, without the benefit of organizational diversity.

A prominent scholar with a keen interest in understanding the economics of information and organization, Arrow (1974), suggests that "...organizations are a means of achieving the benefits of collective action in situations in which the price system fails." Following in a similar vein, it may very well be that R&D communities are a means of achieving the benefits of collective action in situations in which the single organization fails. Given this perspective, it may be helpful to view a community as a meta-organizational structure for accomplishing certain goals reasonably beyond the reach of an individual or single organization.

An organization, such as the firm, can be quite effective in developing a new technology single-handedly, especially if the choice among problem solutions is narrow. The organization facilitates the coordination of researchers' activities, reinforces cooperation, and ensures that the production of information and knowledge is done without duplication and is readily

available to all researchers. However, in situations where the choice among problem solutions is less straightforward, communities may provide a more effective organizing structure. Communities promote a diversity of problem solutions attempted by researchers and enhance competitive pressures to speed the problem solving process. Although communication is more difficult than within a single organization, it can be improved by the grapevine and the institutionalizing mechanisms discussed above.

In the course of this analysis, we have attempted to expose the salient characteristics of communities, all of which are quite familiar phenomena to the researchers who participate in them. Indeed, we have drawn largely on the observations and vocabulary of the researchers themselves in our descriptions of the internal functioning of the community. Many neural network researchers undoubtedly see themselves as members of a larger community, and furthermore, they have an intuitive understanding--indeed, only as they could have--of the dynamics of their field: such notions as bandwagons, critical mass, bootlegging, controversies, grapevines, and momentum are not rarefied abstractions but well-known concepts which are interwoven into the fabric of laboratory life. It is difficult to foresee any comprehensive understanding of emerging technologies that could ignore the role R&D communities might play. The time is ripe for further research.

References

- Allen, Thomas J. (1977) *Managing the Flow of Technology*. Cambridge, MA: MIT Press.
- Allman, William F. (1987) "Designing Computers That Think the Way We Do," *Technology Review*, May/June: 59-65.
- Arrow, Kenneth J. (1974) *The Limits of Organization*. New York: Norton.
- Constant, Edward W. (1980) *The Origins of the Turbojet*. Baltimore: Johns Hopkins University Press.
- DARPA Neural Networks Study*. (1988) Fairfax, VA: AFCEA International Press, 1988.
- Hecht-Nielsen, Robert. (1988) "Neurocomputing: Picking the Human Brain," *IEEE Spectrum* 25, 3: 36-41.
- Hopfield, John J. (1988) "Artificial Neural Networks," *IEEE Circuits and Devices Magazine*, Sept.: 3-10.
- Johnson, George. (1986) *Machinery of the Mind: Inside the New Science of Artificial Intelligence*. Redmond, Wash.: Microsoft Press.
- Katz, Ralph and Thomas J. Allen. (1985) "Organizational Issues in the Introduction of New Technologies," in P.R. Kleindorfer, ed., *The Management of Productivity and Technology in Manufacturing*. New York: Plenum Press.
- Kinoshita, June and Nicholas G. Palevsky. (1987) "Computing with Neural Networks," *High Technology*, May: 24- 31.
- Kohonen, Teuvo. (1988) "An Introduction to Neural Computing," *Neural Networks* 1: 3-16.
- Layton, Edward T. (1974) "Technology as Knowledge," *Technology and Culture* 15: 31-41.
- Malina, Debbie. (1988) "Will the Real Intelligent Machine Please Stand Up? *Mass High Tech* 6, 15: 1, 33.
- National Science Foundation. (1988) *International Science and Technology Data Update: 1988*. (NSF 89- 307) Washington, DC: Government Printing Office.
- Papert, Seymour. (1988) "One AI or Many?" *Daedalus* 117, 1: 1-14.
- Tank, David W. and John J. Hopfield. (1987) "Collective Computation in Neuronlike Circuits," *Scientific American*, December: 104-114.
- von Hippel, Eric A. (1987) "Cooperation Between Rivals: Informal Know-how Trading," *Research Policy* 16: 291-302.
- Widrow, Bernard. (No Date). Letter to member of the International Neural Networks Society.
- Will, Craig A. (1988) "Perceptrons Revisited," *Neural Network Review* 2, 2: 42-49.
- Ziman, John. (1984) *An Introduction to Science Studies*. Cambridge: Cambridge University Press.

Appendix A

PAPERT'S ALLEGORY OF THE TWO SCIENCES: NEURAL NETWORKS AND ARTIFICIAL INTELLIGENCE

Once upon a time two daughter sciences were born to the new science of cybernetics. One sister was natural, with features inherited from the study of the brain, from the way nature does things. The other was artificial, related from the beginning to the use of computers. Each of the sister sciences tried to build models of intelligence, but from very difficult materials. The natural sister built her models (called neural networks) out of mathematically purified neurones. The artificial sister built her models out of computer programs.

In their first bloom of youth the two were equally successful and equally pursued by suitors from other fields of knowledge. They got on very well together. Their relationship changed in the early sixties when a new monarch appeared, one with the largest coffers ever seen in the kingdom of the sciences: Lord DARPA, the Defense Department's Advanced Research Projects Agency. The artificial sister grew jealous and was determined to keep for herself the access to Lord DARPA's research funds. The natural sister would have to be slain.

The bloody work was attempted by two staunch followers of the artificial sister, Marvin Minsky and Seymour Papert, cast in the role of the huntsmen sent to slay Snow White and bring back her heart as proof of the deed. Their weapon was not the dagger but the mightier pen, from which came a book--*Perceptrons*--purporting to prove that neural nets could never fill their promise of building models of mind: *only computer programs could do this*. Victory seemed assured for the artificial sister. And indeed for the next decade all the rewards of the kingdom came to progeny, of which the family of expert systems did the best in fame and fortune.

But Snow White was not dead. What Minsky and Papert had shown the world as proof was not the heart of the princess; it was the heart of a pig. To be more literal: their book was read as proving that the neural net approach to building models of the mind was dead. But a closer look reveals that they demonstrated something less than this. The book did indeed point out very serious limitations of a certain class of nets (nowadays known as one-layer perceptrons) but was misleading in its suggestion that this class of nets was at the heart of connectionism.

Connectionist writings present the story as having a happy ending. The natural sister was quietly nurtured in the laboratories of a few ardent researchers who kept the faith, even when the world at large let itself be convinced that the enterprise was futile.

From, Seymour Papert, "One AI or Many?" *Dædalus*, Volume 117, Number 1, Winter 1988 (p. 4-5)

Date Due

JUN 11 1990		
NOV 11 1991		
MAY 11 1992		
DEC 12 1991		
MAR 04 1993 JUN 08 1993		
MAR 26 1994 OCT 13 2003		
NOV 11 2003		
JAN 05 2004		

MIT LIBRARIES DUPL 1



3 9080 00570434 8

